

WHAT IS CLAIMED IS:

- 1 1. A method for managing files in a file system, comprising:
2 receiving data for a file;
3 storing the data for the file in a plurality of segments;
4 generating an index associated with the file indicating how the file data maps to
5 the segments;
6 receiving an Input/Output request with respect to an address in the file;
7 using the index for the file to determine the segment including data at the
8 requested address in the file; and
9 accessing the determined segment including the data at the requested address.
- 1 2. The method of claim 1, wherein data is stored in the segments by:
2 writing the received file data to one segment; and
3 writing further received data for the file to subsequent segments if the last
4 segment to which the received data was written has no more available space.
- 1 3. The method of claim 1, wherein each segment has a fixed byte length,
2 wherein the index provides a segment order indicating an order in which file data is
3 written to the segments, and wherein the index for the file is used to determine the
4 segment including data at the requested address in the file by:
5 determining an offset into the file including the data at the requested address; and
6 determining an integer quotient value resulting from the offset into the file divided
7 by the fixed byte length, wherein the segment including the data at the requested address
8 is the segment at the integer quotient value in the segment order.
- 1 4. The method of claim 3, further comprising:
2 receiving user input indicating the fixed byte length of each segment.

Sale
9/1
708290" 82446860

1 5. The method of claim 1, further comprising:
2 providing a segment size that is at least greater than a byte size of a largest section
3 within the file; and
4 writing each file section to one segment.

1 6. The method of claim 1, further comprising:
2 storing the segments in a primary storage;
3 copying at least one of the segments in the primary storage onto a secondary
4 storage; and
5 releasing at least one of the segments copied to the secondary storage, wherein
6 space used by the released segment in the primary storage is available for use.

1 7. The method of claim 6, wherein as a result of releasing one or more
2 segments, different segments for one file are capable of being stored in the primary
3 storage and the secondary storage.

1 8. The method of claim 6, wherein accessing the determined segment
2 including the requested address further comprises:
3 determining whether the determined segment is available in the primary storage;
4 and
5 copying the determined segment from the secondary storage to the primary storage
6 if the determined segment is not available in the primary storage.

1 9. The method of claim 6, wherein releasing the segment comprises:
2 storing a partial version of the released segment including less than all data in the
3 segment, wherein the segment data not in the partial version is stored in the secondary
4 storage, wherein the partial version remains on the primary storage after the segment is
5 released.

052301 82445860

1 10. The method of claim 9, wherein the partial version of the determined
2 segment is on the primary storage and wherein accessing the determined segment
3 including the requested address further comprises:
4 accessing the partial version of the determined segment on the primary storage to
5 access the data therein;
6 reaching the end of the partial version when accessing data therein;
7 staging from the secondary storage to the primary storage data from the
8 determined segment that is not in the partial version; and
9 accessing the data from the determined segment staged from the secondary storage
10 to the primary storage.

1 11. The method of claim 9, wherein the partial version is stored only for a
2 first segment of the segments associated with the file.

1 12. The method of claim 6, further comprising:
2 accessing data at the end of the segment, wherein the I/O request requires further
3 file data after accessing the end of the segment;
4 determining from the index a next segment including file data following the file
5 data at the end of the segment data; and
6 accessing the next segment in the primary storage to access the further required
7 file data.

1 13. The method of claim 6, further comprising:
2 maintaining metadata for each segment that is also maintained for files in the file
3 system; and
4 using the metadata for segments and files to determine when to copy segments
5 and files to the secondary storage and when to release segments and files in the primary
6 storage.

108290" 82446860

1 14. The method of claim 13, wherein segments and files in the primary storage
2 are released according to their metadata if used space in the primary storage reaches a
3 threshold level.

1 15. The method of claim 6, wherein the file data in all the segments for the file
2 is capable of being larger than a storage capacity of the primary storage.

1 16. The method of claim 6, further comprising:
2 reading data from one target segment on the secondary storage;
3 determining whether a stage attribute is specified indicating a number of segments
4 to stage ahead; and
5 initiating read requests to stage the number of subsequent segments following the
6 target segment from the secondary storage to the primary storage.

1 17. The method of claim 16, further comprising:
2 receiving user input indicating the number of segments to stage ahead.

1 18. The method of claim 1, wherein the segment does not have a file name and
2 is not represented as a file in the file system.

1 19. The method of claim 1, wherein the index is stored in the file, wherein no
2 user data is stored in the file and all the user data is distributed in the segments.

1 20. A method for managing files in a primary and secondary storage, wherein
2 the secondary storage is comprised of a plurality of drives and storage devices capable of
3 being mounted on the drives, comprising:
4 receiving data for a file;
5 storing the data for the file in a plurality of segments;

03290" 8446860

6 generating an index associated with the file indicating how file data maps to
7 segments; and
8 writing each segment to one of the drives, wherein segments are written to
9 multiple of the drives to distribute the segments across multiple storage devices.

1 21. The method of claim 20, wherein multiple segments are written in parallel
2 to multiple storage devices in multiple drives.

1 22. The method of claim 20, further comprising
2 reading segments on multiple storage devices from multiple drives to stage
3 multiple segments in parallel into the primary storage.

1 23. The method of claim 20, wherein the drives comprise tape drives and
2 wherein the storage devices comprise tape cartridges.

1 24. A system for managing files, comprising:
2 a computer readable medium;
3 a storage system;
4 means for receiving data for a file;
5 means for storing the data for the file in a plurality of segments in the storage
6 device;
7 means for generating an index in the computer readable medium associated with
8 the file indicating how the file data maps to the segments;
9 means for receiving an Input/Output request with respect to an address in the file;
10 means for using the index for the file to determine the segment including data at
11 the requested address in the file; and
12 means for accessing the determined segment including the data at the requested
13 address.

108290" 9244660

1 25. The system of claim 24, wherein the means for storing the for the file in
2 the segments performs:
3 writing the received file data to one segment; and
4 writing further received data for the file to subsequent segments if the last
5 segment to which the received data was written has no more available space.

1 26. The system of claim 24, wherein each segment has a fixed byte length,
2 wherein the index provides a segment order indicating an order in which file data is
3 written to the segments, and wherein means for using the index for the file to determine
4 the segment including data at the requested address in the file performs:
5 determining an offset into the file including the data at the requested address; and
6 determining an integer quotient value resulting from the offset into the file divided
7 by the fixed byte length, wherein the segment including the data at the requested address
8 is the segment at the integer quotient value in the segment order.

1 27. The system of claim 26, further comprising:
2 means for receiving user input indicating the fixed byte length of each segment.

1 28. The system of claim 24, further comprising:
2 means for providing a segment size that is at least greater than a byte size of a
3 largest section within the file; and
4 means for writing each file section to one segment.

1 29. The system of claim 24, wherein the storage system comprises a primary
2 storage, further comprising:
3 a secondary storage;
4 means for copying at least one of the segments in the primary storage onto the
5 secondary storage; and

0954708475860

6 means for releasing at least one of the segments copied to the secondary storage,
7 wherein space used by the released segment in the primary storage is available for use.

1 30. The system of claim 29, wherein as a result of releasing one or more
2 segments, different segments for one file are capable of being stored in the primary
3 storage and the secondary storage.

1 31. The system of claim 29, wherein the means for accessing the determined
2 segment including the requested address further performs:
3 determining whether the determined segment is available in the primary storage;
4 and
5 copying the determined segment from the secondary storage to the primary storage
6 if the determined segment is not available in the primary storage.

1 32. The system of claim 29, wherein the means for releasing the segment
2 performs:
3 storing a partial version of the released segment including less than all data in the
4 segment, wherein the segment data not in the partial version is stored in the secondary
5 storage, wherein the partial version remains on the primary storage after the segment is
6 released.

1 33. The system of claim 32, wherein the partial version of the determined
2 segment is on the primary storage and wherein the means for accessing the determined
3 segment including the requested address further performs:
4 accessing the partial version of the determined segment on the primary storage to
5 access the data therein;
6 reaching the end of the partial version when accessing data therein;

109290" 844686

7 staging from the secondary storage to the primary storage data from the
8 determined segment that is not in the partial version; and
9 accessing the data from the determined segment staged from the secondary storage
10 to the primary storage.

1 34. The system of claim 32, wherein the partial version is stored only for a
2 first segment of the segments associated with the file.

1 35. The system of claim 29, further comprising:
2 means for accessing data at the end of the segment, wherein the I/O request
3 requires further file data after accessing the end of the segment;
4 means for determining from the index a next segment including file data
5 following the file data at the end of the segment data; and
6 means for accessing the next segment in the primary storage to access the further
7 required file data.

1 36. The system of claim 29, further comprising:
2 means for maintaining metadata for each segment that is also maintained for files
3 in the file system; and
4 means for using the metadata for segments and files to determine when to copy
5 segments and files to the secondary storage and when to release segments and files in the
6 primary storage.

1 37. The system of claim 24, wherein segments and files in the primary storage
2 are released according to their metadata if used space in the primary storage reaches a
3 threshold level.

T08230"8246860

1 38. The system of claim 29, wherein the file data in all the segments for the
2 file is capable of being larger than a storage capacity of the primary storage.

1 39. The system of claim 29, further comprising:
2 means for reading data from one target segment on the secondary storage;
3 means for determining whether a stage attribute is specified indicating a number
4 of segments to stage ahead; and
5 means for initiating read requests to stage the number of subsequent segments
6 following the target segment from the secondary storage to the primary storage.

1 40. The system of claim 39, further comprising:
2 means for receiving user input indicating the number of segments to stage ahead.

1 41. The system of claim 24, wherein the segment does not have a file name
2 and is not represented as a file in the file system.

1 42. The system of claim 24, wherein the index is stored in the file, wherein no
2 user data is stored in the file and all the user data is distributed in the segments.

1 43. A system method for managing files, comprising:
2 a primary storage;
3 a secondary storage comprised of a plurality of drives and storage devices capable
4 of being mounted on the drives;
5 means for receiving data for a file;
6 means for storing the data for the file in a plurality of segments on the primary
7 storage;
8 means for generating an index associated with the file indicating how file data
9 maps to segments; and

T08230" B446360

1 44. The system of claim 43, wherein multiple segments are written in parallel
2 to multiple storage devices in multiple drives.

1 46. The system of claim 43, wherein the drives comprise tape drives and
2 wherein the storage devices comprise tape cartridges.

1 48. The article of manufacture of claim 47, wherein data is stored in the
2 segments by:
3 writing the received file data to one segment; and
4 writing further received data for the file to subsequent segments if the last
5 segment to which the received data was written has no more available space.

1 49. The article of manufacture of claim 47, wherein each segment has a fixed
2 byte length, wherein the index provides a segment order indicating an order in which file
3 data is written to the segments, and wherein the index for the file is used to determine the
4 segment including data at the requested address in the file by:

5 determining an offset into the file including the data at the requested address; and

6 determining an integer quotient value resulting from the offset into the file divided
7 by the fixed byte length, wherein the segment including the data at the requested address
8 is the segment at the integer quotient value in the segment order.

1 50. The article of manufacture of claim 49, further comprising:
2 receiving user input indicating the fixed byte length of each segment.

1 51. The article of manufacture of claim 47, further comprising:
2 providing a segment size that is at least greater than a byte size of a largest section
3 within the file; and
4 writing each file section to one segment.

1 52. The article of manufacture of claim 47, further comprising:
2 storing the segments in a primary storage;
3 copying at least one of the segments in the primary storage onto a secondary
4 storage; and
5 releasing at least one of the segments copied to the secondary storage, wherein
6 space used by the released segment in the primary storage is available for use.

1 53. The article of manufacture of claim 52, wherein as a result of releasing one
2 or more segments, different segments for one file are capable of being stored in the
3 primary storage and the secondary storage.

100290" 8446860

1 54. The article of manufacture of claim 52, wherein accessing the determined
2 segment including the requested address further comprises:
3 determining whether the determined segment is available in the primary storage;
4 and
5 copying the determined segment from the secondary storage to the primary storage
6 if the determined segment is not available in the primary storage.

1 55. The article of manufacture of claim 52, wherein releasing the segment
2 comprises:
3 storing a partial version of the released segment including less than all data in the
4 segment, wherein the segment data not in the partial version is stored in the secondary
5 storage, wherein the partial version remains on the primary storage after the segment is
6 released.

1 56. The article of manufacture of claim 55, wherein the partial version of the
2 determined segment is on the primary storage and wherein accessing the determined
3 segment including the requested address further comprises:
4 accessing the partial version of the determined segment on the primary storage to
5 access the data therein;
6 reaching the end of the partial version when accessing data therein;
7 staging from the secondary storage to the primary storage data from the
8 determined segment that is not in the partial version; and
9 accessing the data from the determined segment staged from the secondary storage
10 to the primary storage.

1 57. The article of manufacture of claim 55, wherein the partial version is
2 stored only for a first segment of the segments associated with the file.

09894478-062804

1 58. The article of manufacture of claim 52, further comprising:
2 accessing data at the end of the segment, wherein the I/O request requires further
3 file data after accessing the end of the segment;
4 determining from the index a next segment including file data following the file
5 data at the end of the segment data; and
6 accessing the next segment in the primary storage to access the further required
7 file data.

1 59. The article of manufacture of claim 52, further comprising:
2 maintaining metadata for each segment that is also maintained for files in the file
3 system; and
4 using the metadata for segments and files to determine when to copy segments
5 and files to the secondary storage and when to release segments and files in the primary
6 storage.

1 60. The article of manufacture of claim 59, wherein segments and files in the
2 primary storage are released according to their metadata if used space in the primary
3 storage reaches a threshold level.

1 61. The article of manufacture of claim 52, wherein the file data in all the
2 segments for the file is capable of being larger than a storage capacity of the primary
3 storage.

1 62. The article of manufacture of claim 52, further comprising:
2 reading data from one target segment on the secondary storage;
3 determining whether a stage attribute is specified indicating a number of segments
4 to stage ahead; and

FOIA b 7 - D

5 initiating read requests to stage the number of subsequent segments following the
6 target segment from the secondary storage to the primary storage.

1 63. The article of manufacture of claim 62, further comprising:
2 receiving user input indicating the number of segments to stage ahead.

1 64. The article of manufacture of claim 47, wherein the segment does not have
2 a file name and is not represented as a file in the file system.

1 65. The article of manufacture of claim 47, wherein the index is stored in the
2 file, wherein no user data is stored in the file and all the user data is distributed in the
3 segments.

1 66. An article of manufacture for managing files in a primary and secondary
2 storage, wherein the secondary storage is comprised of a plurality of drives and storage
3 devices capable of being mounted on the drives, by:
4 receiving data for a file;
5 storing the data for the file in a plurality of segments;
6 generating an index associated with the file indicating how file data maps to
7 segments; and
8 writing each segment to one of the drives, wherein segments are written to
9 multiple of the drives to distribute the segments across multiple storage devices.

1 67. The article of manufacture of claim 66, wherein multiple segments are
2 written in parallel to multiple storage devices in multiple drives.

1 68. The article of manufacture of claim 66, further comprising

09894478 "052301
T092290" 8446860

- 2 reading segments on multiple storage devices from multiple drives to stage
- 3 multiple segments in parallel into the primary storage.

- 1 69. The article of manufacture of claim 66, wherein the drives comprise tape
- 2 drives and wherein the storage devices comprise tape cartridges.

090443-05801
T08290" 82445860